Remarks

Applicants respectfully request that this Amendment After Final Action be admitted under 37 C.F.R. § 1.116.

Applicants submit that this Amendment presents claims in better form for consideration on appeal. Furthermore, applicant believes that consideration of this Amendment could lead to favorable action that would remove one or more issues for appeal.

Claims 1, 15, 24 and 31 have been amended. No claims have been canceled. Therefore, claims 1-33 are now presented for examination.

Claims 1-33 stand rejected under 35 U.S.C. §112, second paragraph, as being indefinite for failing to particularly point out and distinctly claim the subject matter which applicant regards as the invention. Applicants submit that claims 1, 15, 24 and 31 have been amended to appear in condition for allowance.

Claims 1-3, 8, 10, 14-17, 21, 23-25 and 31-33 stand rejected under 35 U.S.C. §103(a) as being unpatentable over Shimoi et al. (U.S. Patent No. 5,652,857), Corcoran et al. (U.S. Patent No. 6,449,689), and Naffziger et al. (U.S. Pub. No. 2003/0135694). Applicants submit that the present claims are patentable over any combination of Shimoi, Corcoran and Naffziger.

Shimoi discloses a disk control apparatus for recording and reproducing compression data to physical device of direct access type. The apparatus includes a cache memory between a host computer and a disk drive. The cache memory is divided into a non-compression cache memory for storing non-compression data on a logic block unit basis and a compression cache memory for storing compression data on a

compression group unit basis having the same size as that of the logic sector of the disk drive. A compressing circuit extracts the data stored in the non-compression cache memory on a logic block unit basis and compresses the data. A compression group forming unit collects the compression data of the logic block unit by the compressing circuit unit, thereby forming a compression group and storing the compression group into the compression cache memory. An expanding circuit unit extracts the data stored in the compression cache memory on a compression group unit basis, expands, and develops into the non-compression cache memory. See Shimoi at col. 3, ll. 23-65.

Corcoran discloses a system and method for organizing compressed data on a storage disk to increase storage density. The method and system include a compressor for compressing a data block into a compressed data block, wherein N represents a compression ratio. The storage disk includes a first storage partition having N slots for storing compressed data, and a second storage partition also having N slots for storing overflow data. Each of the N slots in the first partition includes at least one address pointer for pointing to locations in the second partition. According to a further aspect of the system and method, if the compressed data block is less than or equal to 1/N of the data block size, then the compressed data block is stored in a first slot in the first storage partition. If the compressed data block is greater than 1/N of the data block size, then the first 1/N of the compressed data block is stored in the first slot in the first storage partition and a remainder of the compressed block is stored in one or more slots in the second storage partition. The address pointer in the first slot is then updated to point to the one or more slots in the second storage partition. See Corcoran at Abstract.

Naffziger discloses a compression engine for a cache memory subsystem has a pointer into cache tag memory and cache data memory and an interface coupled to the pointer and capable of being coupled to cache tag memory, and cache data memory. The interface reads tag information and uncompressed data from the cache and writes modified tag information and compressed data to the cache. The compression engine also has compression logic for generating compressed data and generate compression successful information, and tag line update circuitry for generating modified tag information according to the compression successful information and the tag information. Also disclosed is a cache subsystem for a computer system embodying the compression engine, and a method of compressing cache using the compression engine. See Naffziger at Abstract.

Claim 1 of the present application recites:

> A computer system comprising:
> a central processing unit (CPU); and
> a cache memory, coupled to the CPU, including:
> a main cache having a plurality of cache lines, each of the plurality of cache lines being compressible to form compressed cache lines to store additional data; and
> a plurality of storage pools to hold a segment of the additional data for a compressed cache line; and
> a cache controller having compression logic to form the compressed cache line by combining a retrieved cache line having a first address comprising a first companion bit value with a companion cache line having a second address comprising a second companion bit value if the companion cache line is resident in the cache memory, wherein the second address differs from the first address by the second companion bit.

Applicants submit that Shimoi, Corcoran and Naffziger all fail to disclose or suggest combining a retrieved cache line having a first address comprising a first companion bit value with a companion cache line having a second address comprising a second companion bit value if the companion cache line is resident in the cache memory. In fact, the Final Office Action acknowledges that Shimoi and Corcoran do not disclose or suggest such a feature. See Final Office Action at page 4, last full paragraph.

However, Naffziger at paragraphs 0056-0060 has been cited as disclosing the feature. Id. at page 5, first full paragraph. Paragraphs 0056-0060 discloses the following:

> [0056] In second alternative embodiment, processor references enter the cache subsystem upon misses in lower-level cache through a processor port 300 (FIG. 3). Each reference address is divided into three portions, a tag address part, a high address bits part, and a location-in-line line part. The tag address part of the reference address is used to find a corresponding tag line in tag memory 304.
>
> [0057] Each line of tag memory 304 in this embodiment has sufficient address tags 400 (FIG. 4) for sixteen-way associativity. Each address tag 400 is associated with a way indicator 402 and flags 404. Among the flags 404 for each address tag is a compressed flag 406, a width indicator 408, and a valid flag 410.
>
> [0058] The associated cache data memory 306 (FIG. 3) stores twelve ways of uncompressed information. At each a location in cache data memory corresponding to a tag memory line is stored a cache line group. The address tags provided in excess of the number required to point to uncompressed data in each cache line group are herein referred to as excess address tags.
>
> [0059] The way indicator 402, 422 (FIG. 4) associated with each address tag indicates where in the cache line group 440 in cache data memory 310 there is stored a cache line associated with the address tag, such as cache lines 412, 432. The width indicator 408, 428

> indicates the width of the cache line as an integral number of sublines.
>
> [0060] The compression engine 308 (FIG. 3) periodically reads the tag memory 304, with its flags 306, and the cache line group 440 from cache data memory 310. The compression engine 308 compresses the cache line group 440 if it is compressible, then rewrites it in compressed form into cache data memory 310. Compression engine 308 also modifies the width indicator 408, way indicator 404, and compressed flags 404 of tag memory flags 306 to correspond with the compressed cache line group 440.

Applicants respectfully submit that nowhere in the above passage is there discloses or suggested a process of combining a retrieved cache line having a first address comprising a first companion bit value with a companion cache line having a second address comprising a second companion bit value if the companion cache line is resident in the cache memory.

Because, Shimoi, Corcoran and Naffziger all fail to disclose or suggest combining a retrieved cache line having a first address comprising a first companion bit value with a companion cache line having a second address comprising a second companion bit value if the companion cache line is resident in the cache memory, any combination of Shimoi, Corcoran and Naffziger would also fail to disclose or suggest such a feature. Therefore, claim 1 is patentable over the combination of Shimoi, Corcoran and Naffziger since neither reference discloses or suggests a process of compressing cache lines.

Claims 2-14 depend from claim 1 and include additional features. Thus, claims 2-14 are also patentable over the combination of Shimoi, Corcoran and Naffziger.

Claim 15 recites:

> A cache memory comprising:

main cache having a plurality of cache lines, each of the plurality of cache lines being compressible to form compressed cache lines to store additional data, wherein a cache line is compressed by combining a retrieved cache line having a first address comprising a first companion bit value with a companion cache line having a second address comprising a second companion bit value if the companion cache line is resident in the cache memory, wherein the second address differs from the first address by the second companion bit; and

a plurality of storage pools to hold a segment of the additional data for a compressed cache line.

For the reasons discussed above with respect to claim 1, claim 15 is also patentable over the combination of Shimoi, Corcoran and Naffziger. Since claims 16-23 depend from claim 15 and include additional features, claims 16-23 are also patentable over the combination of Shimoi, Corcoran and Naffziger.

Claim 24 recites:

A method comprising:

compressing one or more of a plurality of cache lines to form one or more compressed cache lines to store additional data by:

combining a retrieved cache line having a first address comprising a first companion bit value with a companion cache line having a second address comprising a second companion bit value if the companion cache line is resident in the cache memory, wherein the second address differs from the first address by the second companion bit; and

storing a component of the data in one or more of a plurality of storage pools.

Thus, for the reasons discussed above with respect to claim 1, claim 24 is also patentable over the combination of Shimoi, Corcoran and Naffziger. Because claims 25-30 depend from claim 24 and include additional features, claims 25-30 are also patentable over the combination of Shimoi, Corcoran and Naffziger.

Claim 31 recites:

> A computer system comprising:
> a central processing unit (CPU); and
> a cache memory, coupled to the CPU, including:
> main cache having a plurality of cache lines, each of the plurality of cache lines being compressible to form compressed cache lines to store additional data; and
> a plurality of storage pools to hold a segment of the additional data for a compressed cache line; and
> a main memory device coupled to the CPU; and
> a cache controller having compression logic to form the compressed cache line by combining a retrieved cache line having a first address comprising a first companion bit value with a companion cache line having a second address comprising a second companion bit value if the companion cache line is resident in the cache memory, wherein the second address differs from the first address by the second companion bit.

Thus, for the reasons discussed above with respect to claim 1, claim 31 is also patentable over the combination of Shimoi, Corcoran and Naffziger. Because claims 32 and 33 depend from claim 31 and include additional features, claims 32 and 33 are also patentable over the combination of Shimoi, Corcoran and Naffziger.

Claims 4-7, 9, 18-20, 22, and 26-30 stand rejected under 35 U.S.C. §103(a) as being unpatentable over Shimoi et al., Corcoran et al., and Naffziger et al. as applied to claims 3, 8, 17, 21, and 25, respectively, and further in view of Obara (U.S. Patent No. 6,115,787). Applicants submit that the present claims are patentable over any combination of Shimoi, Corcoran, Naffziger and Obara.

Obara discloses storing compressed records into a cache memory of a disk storage system in an easy-to-read manner. Data to be stored in the cache memory is divided into plural data blocks each having two cache blocks in association with track blocks to which the data belongs and are compressed. The respective data blocks after the compression

are stored in one or plural cache blocks. Information for retrieving each cache block from an in-track address for the data block is stored as part of retrieval information for the cache memory. When the respective data blocks in a record is read, the cache block storing the compressed data block is determined based on the in-track address of the data block and the retrieval information. See Obara at Abstract.

However, Obara does not disclose or suggest combining a retrieved cache line having a first address comprising a first companion bit value with a companion cache line having a second address comprising a second companion bit value if the companion cache line is resident in the cache memory. As discussed above, Shimoi, Corcoran and Naffziger all fail to disclose or suggest such a feature. Therefore, any combination of Shimoi, Corcoran, Naffziger and Obara would also not disclose or suggest the feature. As a result, the present claims are patentable over the combination of Shimoi, Corcoran, Naffziger and Obara.

Claims 11-13 stand rejected under 35 U.S.C. §103(a) as being unpatentable over Shimoi et al., Corcoran et al., and Naffziger et al. as applied to claim 10 above, and further in view of Cypher (U.S. Patent No. 6,629,205). Applicants submit that the present claims are patentable over any combination of Shimoi, Corcoran, Naffziger and Cypher.

Cypher discloses a cache memory including a plurality of memory chips that are configured to collectively store a plurality of cache lines. Each cache line includes data and an associated cache tag. The cache tag may include an address tag which identifies the line as well as state information indicating the coherency state for the line. Each cache line is stored across the memory chips in a row formed by corresponding entries

(i.e., entries accessed using the same index address). The plurality of cache lines is grouped into separate subsets based on index addresses, thereby forming several separate classes of cache lines. The cache tags associated with cache lines of different classes are stored in different memory chips. During operation, the cache controller may receive multiple snoop requests corresponding to, for example, transactions initiated by various processors. The cache controller is configured to concurrently access the cache tags of multiple lines in response to the snoop requests if the lines correspond to differing classes. See Cypher at Abstract.

Nonetheless, Cypher does not disclose or suggest combining a retrieved cache line having a first address comprising a first companion bit value with a companion cache line having a second address comprising a second companion bit value if the companion cache line is resident in the cache memory. As discussed above, Shimoi, Corcoran and Naffziger fail to disclose or suggest such a limitation. Therefore, any combination of Shimoi, Corcoran, Naffziger and Cypher would also not disclose or suggest the limitation. Accordingly, the present claims are patentable over the combination of Shimoi, Corcoran, Naffziger and Cypher.

Applicants respectfully submit that the rejections have been overcome, and that the claims are in condition for allowance. Accordingly, applicants respectfully request the rejections be withdrawn and the claims be allowed.

The Examiner is requested to call the undersigned at (303) 740-1980 if there remains any issue with allowance of the case.

Please charge any shortage to our Deposit Account No. 02-2666.

Respectfully submitted,
BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP

Date: 1/22/07

Mark L. Watson
Reg. No. 46,322

12400 Wilshire Boulevard
7th Floor
Los Angeles, California 90025-1026
(303) 740-1980